



Contribution ID: 29

Type: Poster

## Shared I/O components for the ATLAS multi-processing framework

*Thursday, 24 August 2017 16:25 (20 minutes)*

ATLAS uses its multi-processing framework AthenaMP for an increasing number of workflows, including simulation, reconstruction and event data filtering (derivation). After serial initialization, AthenaMP forks worker processes that then process events in parallel, with each worker reading data individually and producing its own output. This mode, however, has inefficiencies: 1) The worker no longer reads events sequentially, which negatively affects data caching strategies at the storage backend. 2) For its non-RAW data ATLAS uses ROOT and compresses across 10-100 events. Workers will only need a subsample of these events, but have to read and decompress the complete buffers. 3) Output files from the individual workers need to be merged in a separate, serial process. 4) Propagating metadata describing the complete event sample through several workers is nontrivial.

To address these shortcomings, ATLAS has developed shared reader and writer components presented in this paper. With the shared reader, a single process reads the data and provides objects to the workers on demand via shared memory. The shared writer uses the same mechanism to collect output objects from the workers and write them to disk. Disk I/O and compression / decompression of data are therefore localized only in these components, event access (by the shared reader) remains sequential and a single output file is produced without merging. Still for object data, which can only be passed between processes as serialized buffers, the efficiency gains depend upon the storage backend functionality.

**Primary authors:** VAN GEMMEREN, Peter (Argonne National Laboratory (US)); MALON, David (Argonne National Laboratory (US)); NOWAK, Marcin (Brookhaven National Laboratory (US)); TSULAIA, Vakho (Lawrence Berkeley National Lab. (US))

**Presenter:** TSULAIA, Vakho (Lawrence Berkeley National Lab. (US))

**Session Classification:** Poster Session

**Track Classification:** Track 1: Computing Technology for Physics Research